

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information, and comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED R&D Status Report 1/1/96-3/31/96	
4. TITLE AND SUBTITLE Applications of the Theory of Distributed and Real Time Systems to the Development of Large-Scale Timing Based Systems				5. FUNDING NUMBERS C - F19628-95-C-0118	
6. AUTHOR(S) Nancy Lynch					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology 77 Massachusetts Avenue Cambridge, MA 02138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Department of the Airforce Electronic Systems Center (AFMC) Hanscom Air Force Base, MA 01731				10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES N/A					
12a. DISTRIBUTION / AVAILABILITY STATEMENT No limits imposed on disclosure <u>DISTRIBUTION STATEMENT A</u> Approved for public release Distribution Unlimited				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words)					
<div style="font-size: 2em; font-weight: bold; margin-bottom: 10px;">19960510 079</div>					
14. SUBJECT TERMS				15. NUMBER OF PAGES 8	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified		18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified		19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	
				20. LIMITATION OF ABSTRACT UL	

NSN 7540-01-280-5500

DRUG QUALITY INSPECTED 1

Standard Form 298 (Rev. 1-39)
Prescribed by ANSI Std. Z39-18
298-102

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to *stay within the lines* to meet optical scanning requirements.

Block 1. Agency Use Only (Leave blank).

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C - Contract	PR - Project
G - Grant	TA - Task
PE - Program Element	WU - Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. (If known)

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD - See DoDD 5230.24, "Distribution Statements on Technical Documents."
DOE - See authorities.
NASA - See Handbook NHB 2200.2.
NTIS - Leave blank.

Block 12b. Distribution Code.

DOD - Leave blank.
DOE - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.
NASA - Leave blank.
NTIS - Leave blank.

Block 13. Abstract. Include a brief (*Maximum 200 words*) factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code (*NTIS only*).

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.



Laboratory for Computer Science

M.L. Dertouzos, Director

R.L. Rivest and A. Vezza, Associate Directors

545 Technology Square, Cambridge, Massachusetts 02139
(617) 253-

April 30, 1996

Mr. Harry Koch
ESC/ENS
5 Eglin Street, Building 1704
Hanscom Airforce Base, MA 01731-2116

Dear Mr. Koch:

This letter contains our R & D Status Report covering the period from Jan. 1, 1996 to Mar. 31, 1996 for Contract F19628-95-C-0118, entitled "Applications of the Theory of Distributed and Real-Time Systems to the Development of Large-Scale Timing-Based Systems".

Technical Progress

Members of MIT's Theory of Distributed Systems group have continued their work on modelling, designing, verifying and analyzing distributed and real-time systems. The focus is on the study of "building-blocks" for the construction of reliable and efficient systems. Our work falls into three general categories: modelling and verification tools, algorithms and impossibility results, and applications. Here, we describe the progress briefly. Sources for more details are provided in the Technical Report.

I. Modelling and verification tools

- Lynch, Segala, Vaandrager and Weinberg have continued working on their new hybrid I/O automaton (HIOA) model, a mathematical model based on labelled transition systems, designed for reasoning about hybrid (continuous/discrete) systems. This quarter, the focus has been on "retrofitting" work done previously using less powerful timed automaton models, so that it rests on the new HIOA model. We have done this with three sets of results from our hybrid systems application project. In addition, work has continued on a full (TR and journal) version of the basic HIOA model paper.
- Garland and Lynch have completed a preliminary design of a Larch interface language for I/O automata. This language, together with associated tools, will facilitate reasoning about I/O automata using the Larch Prover. The language is also intended for use with other tools, such as a simulator and a model-checker.
- Garland and Petrov have polished and revised Petrov's machine verification, using the Larch Prover, of the concurrent timestamp system of Dolev and Shavit.

- Segala has continued work on adapting random walk theory for reasoning about randomized distributed algorithms. He has also developed some new proof rules for analysis using time-like complexity measures, for probabilistic systems. (Segala and Lynch are attempting to finish up the PhD project of Anna Pogoyants, who was killed in a car crash in December. This is part of that project.)

II. Algorithms and impossibility results

- Lynch, Shavit, Shvartsman, and Touitou have revised their paper showing that many important classes of the highly concurrent data structures used for counting and load balancing exhibit nearly linearizable behavior for a broad range of parameters. This paper characterizes the linearizability conditions in terms of a parameter that describes a local property of a low level component, and that does not depend on the size of the data structure. Their revised paper will appear in PODC96.
- Shavit has continued his work on design and analysis of efficient concurrent data structures. This work has many pieces, including a project on Diffracting Priority Queues (with Dan Touitou and Asaph Zemach), one on Diffracting Fetch and Add (with Asaph Zemach), one on Reactive Diffracting Trees (with Giovanni Della Libera – described below), and one on Positive Counter Counting Networks and Applications (with Bill Aiello and Maurice Herlihy). In addition, he is attempting to tie the area together in a monograph on Concurrent Data Structures, co-authored with Maurice Herlihy.

In particular, Shavit and Della Libera have designed a new version of diffracting trees, “Reactive Diffracting Trees”, which grow and shrink according to the load on the data structure. They have obtained results showing that reactive diffracting trees do scale properly.

- Shvartsman is continuing the synthesis of the latest results in the area of parallel computation in the presence of failures and delays. A monograph is in preparation with the target completion date this calendar year. It will be published by Kluwer Academic. A journal paper in this area by Buss, Kanellakis, Ragde and Shvartsman appeared in the *Journal of Algorithms* [2]

III. Applications

A. Distributed system building blocks

- Fekete, Gupta, Luchangco, Lynch and Shvartsman have revised their paper on “eventually serializable data services”, including some optimizations of their original algorithm. Their revised paper will appear in PODC96.

- Shvartsman and Oleg Cheiner, an undergraduate research assistant, have begun implementing a prototype distributed algorithm based on the eventually serializable data service implementation described just above. The prototype will be used as a testbed for exploring algorithm optimizations.
- Lynch and Shvartsman have formulated a specification of a general-purpose processor group-oriented communication primitive. They are applying the primitive to obtain new results and extend previous results for distributed algorithms, e.g., replicated read/write memory. A manuscript is in preparation. With De Prisco they are also developing new and efficient algorithms for the *Do-All* problem of performing n tasks using p message passing processors, while maintaining message and work efficiency. Another manuscript is in preparation.
- Vaziri has nearly completed a proof of correctness of the main algorithm for the main RAID level 5 system. This proof includes much reusable structure, including *recoverability* conditions for the operation graphs used in the algorithm.
- De Prisco has continued his work on modelling, improving, and verifying the practical Paxos algorithm for fault-tolerant distributed consensus. The proof of the main algorithm is nearly finished, but there is work still to be done in the subsidiary algorithms (for leader election and failure detection), as well as in applications of the Paxos algorithm to replicated data management and distributed transaction processing.

B. Transit

- Weinberg and Lynch have completed their analysis, using hybrid I/O automata, invariants and simulation mappings, of a collection of typical vehicle deceleration maneuvers. These appear in Weinberg's M.S. thesis.
- Lynch, Weinberg and Delisle have completed a paper on modelling and analyzing separate vehicle protection (VP) subsystems, as used in the Raytheon Personal Rapid Transit project, for the proceedings of the DIMACS-95 Workshop on Hybrid Systems. Lynch has begun working with M.Eng. student Carl Livadas on extensions of this work to handle more types of safety subsystems.
- Lynch has completed her work on the analysis of an acceleration maneuver, using three levels of abstraction. She has written a final report for the AMAST Workshop on Hybrid Systems.
- Lynch has begun working with undergrad student Kate Dolgin and postdoc Mike Branicky on modular safety analysis for the platoon join maneuver of the California PATH intelligent highway project.

C. Communication

- Smith has finished his formal verification of TCP. After finishing this, he started his verification of T/TCP, using a simulation mapping to TCP, and discovered that it does not implement TCP! He discovered a situation in which T/TCP delivers duplicate data to the user. He will proceed on this project by determining and proving the weaker properties that T/TCP does in fact guarantee, and by considering whether the stronger properties are actually possible to achieve. He will seek either an improved model or an impossibility result.

D. Probabilistic Systems

- Lynch and Segala are working intensively to complete the work by Pogonyants and Segala on modelling and proof of the (randomized) Aspnes-Herlihy consensus protocol. The safety proof has been completed. There is a good draft of the proof of the probabilistic progress properties, but this needs more work.

Special Programs and Major Items of Equipment

None.

Changes in Key Personnel

1. M.S. student H.B. Weinberg has finished his M.S. thesis work and has left M.I.T.
2. M.Eng. student Carl Livadas has joined the group to complete the work on real-time-systems modelling and verification begun by H.B. Weinberg.
3. Undergraduate student Kate Dolgin has joined the group to help with the transit modelling project.

Trips, Talks and Conferences

1. Two meetings were held involving members of our group and Lincoln Labs researchers involved in evaluating air traffic control and aircraft control systems. These were on Jan. 19 and Feb. 23.
2. Lynch gave two invited addresses at the University of Florida, in Gainesville, one on distributed shared memory and multicast, and the other on modelling automated transit systems. Both were on Jan. 26.

3. Roberto Segala visited from Bologna to work with Lynch on probabilistic system verification, for two weeks in February.
4. Lynch attended the annual ARPA Networking P.I. meeting in Charleston, S.C., organized by Gary Minden, on Feb. 25-27. She spoke about the general enterprise of defining/studying building blocks for high-assurance distributed systems, and about specific results on distributed shared memory and multicast, on eventually serializable data services, on TCP and T/TCP communication services, and on counting networks and other efficient concurrent data structures.
5. Lynch attended the AMAST Workshop on Real-Time Systems in Salt Lake City, Mar. 4-5. She gave an invited address on modelling and verification of transit systems.
6. Lynch and Livadas visited Raytheon in to discuss current designs of automated transit systems, on Mar. 21.
7. Mark Smith presented his paper entitled "Formal Verification of TCP" at The Second Technical Conference on Telecommunications R&D in Massachusetts, Lowell, MA, in March.
8. Shvartsman participated in the Information Survivability-Formal Methods PI Conference on January 16-18, 1996 in San Diego, CA. He gave a talk on research direction in TDS titled "Theory of Distributed and Real-Time Systems".

Areas of Concern

None.

Statement of Sufficiency

The contractually prescribed effort appears to be sufficient to achieve the objectives of this contract.

Degrees awarded

1. H.B. Weinberg, M.Eng. thesis entitled "Correctness of Vehicle Control Systems - A Case Study". Completed in Jan., 1996.

Related Accomplishments

During this reporting period, the following papers have been submitted to, been accepted to, or have appeared in journals and conference proceedings:

References

- [1] Hagit Attiya, Amir Herzberg, and Sergio Rajsbaum. Optimal clock synchronization under different delay assumptions. *SIAM Journal on Computing*, April 1996.
- [2] J.F. Buss, P.C. Kanellakis, P. L. Ragde, and A. Shvartsman. Parallel algorithms with processor failures and delays. *Journal of Algorithms*, 20:45–86, January 1996.
- [3] Alan Fekete, David Gupta, Victor Luchangco, Nancy Lynch, and Alex Shvartsman. Eventually-serializable data services. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, Philadelphia, PA, May 1996. To appear.
- [4] Nancy Lynch. A three-level analysis of a simple acceleration maneuver, with uncertainties. In *Proceedings of the Third AMAST Workshop on Real-Time Systems*, pages 1–22, Salt Lake City, Utah, March 1996.
- [5] Nancy Lynch and Sergio Rajsbaum. On the Borowsky-Gafni simulation algorithm. In *Proceedings of ISTCS 1996: The Fourth Israel Symposium on Theory of Computing and Systems*, Jerusalem, Israel, June 1996. To appear. Also, short version to appear in *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, Philadelphia, PA, May 1996.
- [6] Nancy Lynch, Nir Shavit, Alex Shvartsman, and Dan Touitou. Counting networks are practically linearizable. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, Philadelphia, PA, May 1996. To appear.
- [7] Tsvetomir P. Petrov, Anna Pogoyants, Stephen J. Garland, and Nancy A. Lynch. Computer-assisted verification of an algorithm for concurrent timestamps, April 1996. Submitted for publication.
- [8] Nir Shavit and Dan Touitou. Elimination trees and the construction of pools and stacks, February 1996. Submitted for publication.
- [9] Nir Shavit and Dan Touitou. Software transactional memory, February 1996. Submitted for publication.
- [10] Nir Shavit and Asap Zemach. Diffracting trees. Submitted for publication. Also, in *Proceedings of the Annual Symposium on Parallel Algorithms and Architectures (SPAA)*, June 1994.
- [11] Mark Smith. Formal verification of communication protocols, April 1996. Submitted for publication.

- [12] Mark Smith. Formal verification of TCP. In *Proceedings of The Second Technical Conference on Telecommunications R&D in Massachusetts*, pages 279–299, Lowell, MA, March 1996.
- [13] H.B. Weinberg, Nancy Lynch, and Norman Delisle. Verification of automated vehicle protection systems. In *DIMACS Workshop on Verification and Control of Hybrid Systems*, October 1995. To appear in R. Alur, T. Henzinger, and E. Sontag, editors, *Hybrid Systems III*, Lecture Notes in Computer Science, Springer-Verlag.

Also, Lynch's book on Distributed Algorithms appeared in print during this reporting period:

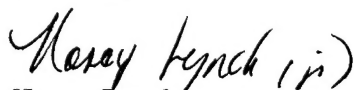
Nancy Lynch.

Distributed Algorithms.

Morgan Kaufmann Publishers, Inc., San Mateo, CA, March 1996.

Award: Undergraduate student Tsvetomir Petrov became the recipient of the first annual Anya Pogogyants Undergraduate Research Award.

Sincerely,



Nancy Lynch

Cecil H. Green Professor

Electrical Engineering and Computer Science

(617)253-7225

lynch@theory.lcs.mit.edu

MIT Laboratory for Computer Science
Applications of the Theory of Distributed Real-Time Systems
To the Development of Large-Scale Timing-Based Systems
Prof. Nancy Lynch, Principal Investigator

R & D Status Report
 Program Financial Status
 ARPA Contract # F19628-95-C-0118
 CLIN # 0002
 1996 First Quarter (1/96 - 3/96)

Total Base Contract
 Current Funding Profile
 Equipment

Planned Expenditures	Actual Expenditures at Report Date	% Completion	Budget At Completion	Latest Revised Estimate	Remarks
858,443	123,334	14.37%	858,443	858,443	
363,787	123,334	33.90%		123,334	*
35,308	0	0.00%			**

* Data reflects all received funding. Current funding is sufficient for this fiscal year. Next fiscal year's anticipated funding requirements are \$188,800

** Equipment funding is for 3 budgeted workstations. None have been purchased as yet.